



<http://www.diva-portal.org>

Preprint

This is the submitted version of a paper presented at *International Conference on Medical Image Computing and Computer Assisted Intervention, MICCAI 2024, 27th International Conference, Marrakesh, Morocco, October 6–10, 2024*.

Citation for the original published paper:

Gunnarsson, N., Sjölund, J., Kimstrand, P., Schön, T B. (2024)
Online Learning in Motion Modeling for Intra-interventional Image Sequences
In: Marius George Linguraru; Qi Dou; Aasa Feragen; Stamatia Giannarou; Ben Glocker; Karim Lekadir; Julia A. Schnabel (ed.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024: 27th International Conference, Marrakesh, Morocco, October 6–10, 2024, Proceedings, Part II* (pp. 706-716). Springer
Lecture Notes in Computer Science (LNCS)
https://doi.org/10.1007/978-3-031-72069-7_66

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-538033>

Online learning in motion modeling for intra-interventional image sequences

Niklas Gunnarsson^{1,2}[0000–0002–9013–949X], Jens Sjölund¹[0000–0002–9099–3522],
Peter Kimstrand²[0000–0001–9667–5595], and
Thomas B. Schön¹[0000–0001–5183–234X]

¹ Department of Information Technology, Uppsala University, Sweden
`{firstname}.{surname}@it.uu.se`

² Elekta Instrument AB, Stockholm, Sweden
`{firstname}.{surname}@elekta.com`

Abstract. Image monitoring and guidance during medical examinations can aid both diagnosis and treatment. However, the sampling frequency is often too low, which creates a need to estimate the missing images. We present a probabilistic motion model for sequential medical images, with the ability to both estimate motion between acquired images and forecast the motion ahead of time. The core is a low-dimensional temporal process based on a linear Gaussian state-space model with analytically tractable solutions for forecasting, simulation, and imputation of missing samples. The results, from two experiments on publicly available cardiac datasets, show reliable motion estimates and an improved forecasting performance using patient-specific adaptation by online learning.

Keywords: Image registration · Online learning · Dynamic probabilistic modeling.

1 Introduction

Sequential imaging during medical interventions, so-called intra-interventional imaging, appears in several medical examinations. In cardiology, diagnostic decisions may be supported by cardiac ultrasound or cardiac MRI by acquiring images of the heart over one or several cardiac cycles [1]. In MR-guided radiotherapy [21], 2D cine MRI is used to monitor moving tumors and organs at risk during ongoing treatment sessions. This enables controlling and adapting the treatment beam [13].

A common desire is to identify the anatomical motion from a static reference image to each subsequent image in the temporal sequence. This enables transferring segmentations (e.g. organs) identified in the reference image and estimating their location in the sequence. Finding the corresponding deformation field is the main goal of motion estimation. Of particular interest are diffeomorphic deformations, which are topology-preserving and ensure one-to-one mapping between the pixels/voxels in the two images. Examples of conventional diffeomorphic image registration methods are Large Deformation Diffeomorphic

Metric Mapping [5] and symmetric normalization [3]. Recently, deep learning image registration methods [15, 25] have shown fast and accurate performance in motion detection and organ tracking by removing the iterative optimization procedure from inference time and solving tasks in nearly real-time. However, image registration methods do not consider the sequential nature of an image sequence and estimate the motion using one image pair at a time.

With sequential images, an interesting research question is to model the temporal sequence from the data. We refer to this as motion modeling – a model with the possibility of estimating the motion at the previous, current, and future times. We present a diffeomorphic motion model suitable for intra-interventional medical image sequences. For this, we define and model a temporal process in a low-dimensional latent space with the possibility to impute and forecast missing samples in the sequence. Furthermore, our model is the first, to the best of our knowledge, to support online learning of the temporal model which makes it suitable for real-time scenarios.

2 Related work

Motion modeling and real-time analysis of intra-interventional medical images is a relatively new research direction. The literature shows that the most common motion to analyze is cyclic patterns like cardiac or respiratory motion [7, 20]. A general approach is to embed the image data into a lower dimensional space and model the temporal process in this domain. Romaguera et al. [24] present a forecasting approach where they suggest a convolutional LSTM to extrapolate the temporal process in the latent dimension. Extension of their work includes a forecasting 4D motion (3D + time) given 2D intra-interventional images using a probabilistic setting [23]. Krebs et al. [14] proposed a more general probabilistic motion model. Their model relies on a conditional variational autoencoder, where they approximate the posterior distribution in the latent dimension using a temporal convolutional neural network. During training, they minimize the KullbackLeibler divergence between their approximate posterior distribution and a known Gaussian process prior. Missing samples in the sequence are then replaced with samples from this prior. Their work shows reliable diffeomorphic estimates of the displacement field with imputing and forecasting possibility. However, they are limited to image sequences of fixed length. To overcome this, Gunnarsson et al. [11] modeled the low dimensional temporal process using a linear Gaussian state space model, i.e., a first-order Markov process. The work we present here is a further development of this model, including support for online learning and architectural improvements.

3 Background - Linear Gaussian state space model

A linear Gaussian state space model (LG-SSM) is a linear representation of a state space model. The model defines a first-order Markov process for a dynamic

state variable $z_t \in \mathbb{R}^p$ followed by a transmission operation between the state variable and an observed variable $x_t \in \mathbb{R}^q$, i.e.

$$z_t|z_{t-1} \sim \mathcal{N}(z_t|Az_{t-1}, Q), \quad x_t|z_t \sim \mathcal{N}(x_t|Cz_t, R), \quad z_0 \sim \mathcal{N}(z_0|\mu_0, P_0), \quad (1)$$

where $A \in \mathbb{R}^{p \times p}$ and $C \in \mathbb{R}^{q \times p}$ denote the state and observation matrix, respectively, $Q \in \mathbb{R}^{p \times p}$ and $R \in \mathbb{R}^{q \times q}$ denote covariance matrices for uncertainties and μ_0, P_0 is the initial values of the state process. Besides that, LG-SSMs are beneficial since the state prediction, $z_{t+k} | z_t$, and smoothing, $z_{t-k} | z_t$, $k > 0$, are analytically tractable using i.e. Kalman filtering [12] and RTS smoothing [22]. Lately, parameter-estimated LG-SSM has shown impressive results in long-range sequence modeling tasks, outperforming recent methods like RNNs, CNNs, and Transformers[10].

To reduce the computational complexity of high-dimensional sequences $\mathbf{y} = [y_1, \dots, y_T]$, like videos, Fraccaro et al. [9] proposed a probabilistic dynamical model that embeds the sequence into a lower-dimensional space where it is represented as an LG-SSM.

4 Method

Given the data $\{(y_0, \mathbf{y})^{(i)}\}_{i=1}^n$ of static reference images y_0 and time sequences $\mathbf{y} = [y_1, \dots, y_T]$ our goal is to first model the spatiotemporal changes and then use this model to reconstruct and generate samples at other times. For this, we explain the spatiotemporal changes as the spatial transformation φ_t from the static reference image to each time step t in the sequence such that $y_t \approx y_0 \circ \varphi_t$ [11]. To include spatial information in the transformation, like contour information and description of shapes, but still limit the temporal process to the most significant temporal changes, we estimate φ_t based on the temporal process and spatial information s given the static reference image y_0 , i.e.,

$$\varphi_t = g_{\theta_g}(x_t, s), \quad s = f_{\theta_s}(y_0), \quad (2)$$

where x_t is a low-dimensional variable at time t . By doing this, we can separate the temporal changes and characteristic features from the images within the sequence. Since the image process may be incomplete due to e.g. missing samples, we define the spatial information based on the static reference image y_0 only. The spatial transformation φ_t is a function of x_t and s (and s is a function of y_0), we can parameterize $p_{\theta}(\mathbf{y} | y_0, \mathbf{x}) = \prod_{t=1}^T p_{\theta}(y_t | y_0, x_t)$ with a generative network with parameters $\theta = \{\theta_g, \theta_s\}$ and model the likelihood $p_{\theta}(y_t | y_0, x_t) = p_{\theta}(y_t | y_0, \varphi_t)$ as any computable continuous distribution in θ . Furthermore, to estimate missing samples in the sequence we model the temporal process in the lower dimension using an LG-SSM, driven by the state variables z_t and with parameters $\gamma = \{A, Q, C, R, \mu_0, \Sigma_0\}$. Finally, given an approximate posterior $q(\mathbf{x}, \mathbf{z} | y_0, \mathbf{y}) = q_{\phi}(\mathbf{x} | y_0, \mathbf{y})p_{\gamma}(\mathbf{z} | \mathbf{x})$ an evidence lower bound (ELBO) can be derived as

$$\log p(\mathbf{y} | y_0) \geq \mathbb{E}_{q_{\phi}(\mathbf{x}|y_0, \mathbf{y})} \left[\log \frac{p_{\theta}(\mathbf{y} | y_0, \boldsymbol{\varphi})}{q_{\phi}(\mathbf{x} | y_0, \mathbf{y})} + \mathbb{E}_{p_{\gamma}(\mathbf{z}|\mathbf{x})} \left[\log \frac{p_{\gamma}(\mathbf{x}, \mathbf{z})}{p_{\gamma}(\mathbf{z} | \mathbf{x})} \right] \right], \quad (3)$$

where $p_\gamma(\mathbf{x}, \mathbf{z})$ and $p_\gamma(\mathbf{z} | \mathbf{x})$ are both analytical tractable using Kalman filtering and RTS smoothing. During the training process we maximize the approximate ELBO by sampling $\tilde{\mathbf{x}} \sim q_\phi(\mathbf{x} | y_0, \mathbf{y})$ and $\tilde{\mathbf{z}} \sim p_\gamma(\mathbf{z} | \tilde{\mathbf{x}})$ and update the parameters of the inference network (ϕ), the LG-SSM (γ) and the generative network (θ) simultaneously. For a complete derivation of the ELBO, we refer to supplemental material. A schematic overview of our probabilistic model given the observed variables \mathbf{y}, y_0 and unobserved variables \mathbf{x}, \mathbf{z} is shown in Fig. 1.




Fig. 1. An observation y_t is downsampled to the temporal process (a). The spatial transformation is generated given y_0 and the low-dimensional motion model (b). (c) visualizes the entire model.

4.1 Online learning

To adapt the model for individual patient motion, we propose a fast online learning procedure that operates on the motion model only. This means we only focus on the LG-SSM parameters γ and keep the inference and generative network parameters $\{\phi, \theta\}$ fixed (shown in Fig. 2). To update the LG-SSM parameters, we iteratively maximize the exact marginal log-likelihood for the N most recent samples of the temporal process at each sampling time t , i.e.,

$$\max_{\gamma_t} \log p_{\gamma_t}(x_{t-N:t}) = \max_{\gamma_t} \log \prod_{k=t-N}^t p_{\gamma_t}(x_k | x_{k-1}). \quad (4)$$

This approach is based on the moving horizon estimation technique [16], which is a well-established method for state estimation in real-time applications. We calculate the marginal log-likelihood using the Kalman filter and update the parameters using gradient-based optimization methods. The algorithm for our proposed online learning procedure is shown in Algorithm 1.

4.2 Implementation details

In our implementation, the inference network and the spatial feature extraction share a similar network architecture. We downsample the data using convolutional layers with filters [32, 32, 32, 16], extract the spatial features at each


Algorithm 1 Online training**Require:** $\gamma_0, \phi, \theta, y_0, N$ $\mathbf{s} \leftarrow f_\phi(y_0)$ $x \leftarrow \{\}$ **while** y_t arrives **do** $\tilde{x}_t \sim q_\theta(x_t | y_0, y_t)$ $x \leftarrow (x \setminus \tilde{x}_{t-N}) \cup \tilde{x}_t$ \triangleright Data collection $p_{\gamma_{t-1}}(x) \leftarrow \text{Kalman Filter}(x)$ $\mathcal{L} \leftarrow \log p_{\gamma_{t-1}}(x)$ \triangleright Equation (4) $\gamma_t \leftarrow \text{update}(\gamma_{t-1}, \nabla_{\gamma_{t-1}} \mathcal{L})$ 

Fig. 2. During online-learning we fix the parameters of the encoder (ϕ) and the decoder (θ), and only update the parameters of the LG-SSM (γ).

level, and estimate the mean and covariance of x_t at the bottom level. For the LG-SSM, we use eight dimensions for x_t ($p = 8$) and 16 for the state-variable z_t ($q = 16$) and estimate all model parameters. In the generative network, we use attention gates [18] to focus the temporal changes on the spatial features of the reference image at each resolution and use the same number of filters as in the inference network. To ensure diffeomorphic estimation of φ_t we consider the output as the stationary velocity field v_t and first smooth it using a Gaussian filter [15] and then compute the transformation numerically using four scaling-and-squaring layers [2], a proven approach to obtain diffeomorphic registrations [8, 15]. Our implementation is publicly available at https://github.com/ngunnar/2D_motion_model and for a more detailed description, we refer to supplemental material.

5 Experiments

For experiments, we evaluate our model on two publicly available datasets: i) single-cycle cine-MRI sequences from the Automatic Cardiac Diagnosis Challenge (ACDC) [6], and ii) longer sequences of cardiac ultrasound images from the EchoNet-Dynamic database [19]. On the EchoNet-Dynamic dataset, we perform online learning with a moving horizon of $N = 75$. This is not suitable for the ACDC dataset as the sequences are too short. Instead, on the ACDC dataset

we evaluate the capability to reconstruct the sequence from sparsely sampled sequences where we only observe a subset of the images in the sequence.

Both datasets are segmented manually at the end-systole and the end-diastole time points. We use the first time point as our static reference image and calculate the Dice score coefficients (DSC) and the 95%-th Hausdorff distance (HD95) between the other manually segmented frame and our estimation at the given point for evaluation. We compare our registration accuracy against no estimated motion and two well-established image registration methods: symmetric normalization (SyN) [3] and elastic registration [17], both using the ANTs software [4]. Moreover, for online learning, we leave a horizon of $H = 50$ samples for each sequence and calculate both the log-likelihood of the unseen sequence $x_{T:T+H}$ and the RMSE between 50 samples from the forecasting distribution and the true latent values. Furthermore, we also calculated the Dice score between the samples 25 steps ahead and the estimated segmentation given the entire sequence. In Table 1, we present the overall result from both datasets. All models produce diffeomorphic deformations (positive Jacobian determinants) and this metric is omitted from the table. The execution times for motion estimation with and without online learning are approximately 15 ms and 75 ms on a single CPU, respectively.




Fig. 3. Overlay of true sequence (magenta), and $\varphi_t = 0$, on top, and our estimation given every 10th sample, on bottom (green). On the right, the distribution of the left ventricle area for 20 latent samples under three scenarios: all time points observed, every 5th, and every 10th. The figure is colored in the online version.

ACDC: In the ACDC experiment, which consists of 100 patients for training and 50 for testing, we resample the images with spacing 1.5×1.5 mm and crop it to 128×128 pixels with the ventricles in the center. The original sequences are in 3D with limited resolution in one orientation. Therefore we only consider the 2D motion in the other two orientations. For training, we split the volume into slices and removed slices with no annotations, resulting in a training set of 840 sequences. In the evaluation part, we use the middle slice of the volume

in the test dataset to ensure connected segmented regions with no mismatch due to out-of-plane motions. For consistency regarding the sequence length, we resample each sequence to 35 samples using bilinear interpolation. During the training phase, we augment the data using random rotation, flip, and translation of the whole sequences. For evaluation, we use the segmented regions of the right ventricle (RV), the left ventricle myocardium (LV-Myo), and the left ventricle blood pool (LV-BP). Fig. 3 shows the result from where we reconstruct the entire sequence using only every 10th sample as input to our model.

EchoNet-Dynamic: The EchoNet-Dynamic dataset includes 10 023 unique cardiac ultrasound videos of various lengths with left ventricle segmentations (LV). We split this data into a training set of 9 540 videos and 483 videos for testing. Furthermore, during training, for each epoch, we randomly selected a sequence of 50 frames from each video. Fig. 4 shows the result of online training when we forecast the motion 50 time-steps ahead.

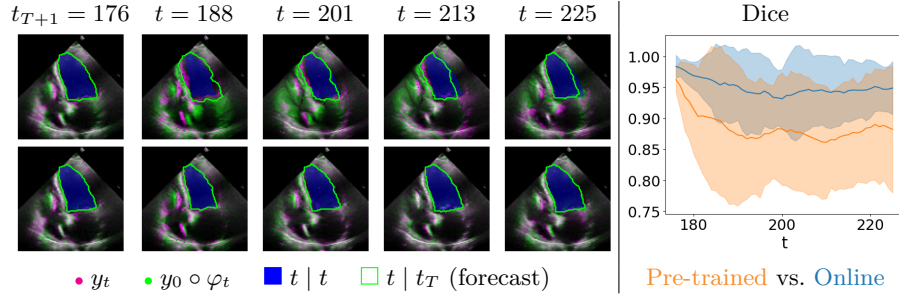


Fig. 4. Overlay between true sequence (magenta) and forecasted sequence (green) using pre-trained model, on top, and online learning, on bottom. To the right, Dice score distribution of the left ventricle from 20 forecasted samples and the estimated region given the entire sequence. The figure is colored in the online version.

6 Discussion and conclusion

In this work, we have presented a motion model for intra-interventional medical images. We define the motion model in a low-dimensional space as a probabilistic LG-SSM with analytical solutions to the inference problem, like imputation for undersampled data (smoothing) and forecasting into the future (prediction). In the first experiment, on the ACDC dataset, we show a marginally improved accuracy compared to well-established diffeomorphic image registration methods, even in cases where we subsample the data and retain only 10% of the original sequence. Our model, operating in a lower and more manageable latent space, shows similar accuracy to recent work [14]. However, a direct comparison is not feasible since both the code and some of the data are not publicly available. In

Table 1. Overall results from the two datasets.

	ACDC						EchoNet	
	DSC			HD95[mm]			DSC	HD95 ¹
	RV	LV-Myo	LV-BP	RV	LV-Myo	LV-BP	LV	LV
None	0.70	0.52	0.66	9.11	8.46	11.06	0.74	8.37
Elastic	0.77	0.72	0.81	6.81	5.88	6.33	0.87	5.94
SyN	0.79	0.72	0.86	5.88	5.22	4.69	0.88	4.68
Our	0.80	0.82	0.84	5.35	4.34	5.18	0.86	4.92
Our 5th	0.80	0.81	0.84	5.28	4.44	5.27	—	—
Our 10th	0.79	0.80	0.83	5.51	4.63	5.57	—	—

	EchoNet Forecasting		
	$\log p_\gamma(x_{T:T+H})$	RMSE ($x_{T:T+H}$)	DSC ($\varphi_{T+25 T}$)
Pre-trained	−10.5	7.04	0.81
Online	−6.3	5.54	0.85

¹ Spacing is not specified in the dataset. The metric is given in pixels.

the second experiment, on the EchoNet-dynamics dataset, we show the capacity of the model to adapt to new, patient-specific data by using online learning and updating the weights in the low-dimensional LG-SSM. The online learned model shows forecasting improvements in both similarities of the latent process with higher likelihood given the true process and lower distance between the samples compared to the pre-trained model as well as the calculated Dice score for predicted samples. The registration accuracy in this experiment is slightly worse than the conventional image registration methods, and can hopefully be improved by refining the hyperparameter settings or data preprocessing. Finally, we believe patient-specific adaptation and reliable forecasting predictions are necessary for longer sequences to support advanced procedures, like real-time adaptation in MR-guided radiotherapy. Other topics for further investigation include relating the uncertainty in the latent temporal process to the uncertainty in the estimated displacement field and observing how each component of the latent space contributes to the actual motion.

Acknowledgments. This research was partially supported by the *Wallenberg AI, Autonomous Systems and Software Program (WASP)* funded by Knut and Alice Wallenberg Foundation.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

Bibliography

- [1] Angelini, E.D., Laine, A.F., Takuma, S., et al.: LV volume quantification via spatiotemporal analysis of real-time 3-D echocardiography. *IEEE Transactions on Medical Imaging* **20**(6), 457–469 (2001)
- [2] Arsigny, V., Commowick, O., Pennec, X., Ayache, N.: A log-euclidean framework for statistics on diffeomorphisms. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI: 9th International Conference, Copenhagen, Denmark, October. Proceedings, Part I* 9. Springer (2006)
- [3] Avants, B.B., Epstein, C.L., Grossman, M., Gee, J.C.: Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. *Medical image analysis* **12**(1), 26–41 (2008)
- [4] Avants, B.B., Tustison, N., Song, G., et al.: Advanced normalization tools (ANTs). *Insight j* **2**(365), 1–35 (2009)
- [5] Beg, M.F., Miller, M.I., Trounev, A., Younes, L.: Computing large deformation metric mappings via geodesic flows of diffeomorphisms. *International journal of computer vision* **61**, 139–157 (2005)
- [6] Bernard, O., Lalande, A., Zotti, C., et al.: Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: is the problem solved? *IEEE Transactions on Medical Imaging* **37**(11), 2514–2525 (2018)
- [7] Chen, C., Qin, C., Qiu, H., et al.: Deep learning for cardiac image segmentation: a review. *Frontiers in Cardiovascular Medicine* **7**, 25 (2020)
- [8] Dalca, A.V., Balakrishnan, G., Guttag, J., Sabuncu, M.R.: Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces. *Medical image analysis* **57**, 226–236 (2019)
- [9] Fraccaro, M., Kamronn, S., Paquet, U., Winther, O.: A disentangled recognition and nonlinear dynamics model for unsupervised learning. *Advances in neural information processing systems* **30** (2017)
- [10] Gu, A., Goel, K., Ré, C.: Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396* (2021)
- [11] Gunnarsson, N., Sjölund, J., Kimstrand, P., Schön, T.B.: Unsupervised dynamic modeling of medical image transformations. In: *2022 25th International Conference on Information Fusion (FUSION)*. pp. 01–07. IEEE (2022)
- [12] Kalman, R.E.: A new approach to linear filtering and prediction problems. *Journal of basic Engineering* **82**(1), 35–45 (1960)
- [13] Keall, P.J., Brighi, C., Glide-Hurst, C., et al.: Integrated MRI-guided radiotherapy opportunities and challenges. *Nature Reviews Clinical Oncology* **19**(7), 458–470 (2022)
- [14] Krebs, J., Delingette, H., Ayache, N., Mansi, T.: Learning a generative motion model from image sequences based on a latent motion matrix. *IEEE Transactions on Medical Imaging* **40**(5), 1405–1416 (2021)

- [15] Krebs, J., Delingette, H., Mailhé, B., et al.: Learning a probabilistic model for diffeomorphic registration. *IEEE Transactions on Medical Imaging* **38**(9), 2165–2176 (2019)
- [16] Mattingley, J., Boyd, S.: Real-Time Convex Optimization in Signal Processing. *IEEE Signal Processing Magazine* **27**(3), 50–61 (2010)
- [17] Modersitzki, J.: Numerical methods for image registration. OUP Oxford (2003)
- [18] Oktay, O., Schlemper, J., Folgoc, L.L., et al.: Attention U-Net: Learning where to look for the pancreas. arXiv preprint arXiv:1804.03999 (2018)
- [19] Ouyang, D., He, B., Ghorbani, A., et al.: Video-based AI for beat-to-beat assessment of cardiac function. *Nature* **580**(7802), 252–256 (2020)
- [20] Paganelli, C., Whelan, B., Peroni, M., et al.: MRI-guidance for motion management in external beam radiotherapy: current status and future challenges. *Physics in Medicine & Biology* **63**(22), 22TR03 (2018)
- [21] Raaymakers, B.W., Lagendijk, J., Overweg, J., et al.: Integrating a 1.5 T MRI scanner with a 6 MV accelerator: proof of concept. *Physics in Medicine & Biology* **54**(12), N229 (2009)
- [22] Rauch, H.E., Tung, F., Striebel, C.T.: Maximum likelihood estimates of linear dynamic systems. *AIAA journal* **3**(8), 1445–1450 (1965)
- [23] Romaguera, L.V., Mezheritsky, T., Mansour, R., et al.: Probabilistic 4D predictive model from in-room surrogates using conditional generative networks for image-guided radiotherapy. *Medical image analysis* **74**, 102250 (2021)
- [24] Romaguera, L.V., Plantefève, R., Romero, F.P., et al.: Prediction of in-plane organ deformation during free-breathing radiotherapy via discriminative spatial transformer networks. *Medical image analysis* **64**, 101754 (2020)
- [25] Ye, M., Yang, D., Huang, Q., et al.: SequenceMorph: A Unified Unsupervised Learning Framework for Motion Tracking on Cardiac Image Sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(08), 10409–10426 (2023)

Supplemental Materials:

Online learning in motion modeling for intra-interventional image sequences

Derivation of the ELBO

The conditional probability density function

$$p(\mathbf{y} \mid y_0) = \frac{p(\mathbf{x}, \mathbf{y}, \mathbf{z} \mid y_0)}{p(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y})}, \quad (1)$$

is infeasible due to the intractable posterior distribution $p(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y})$. Instead, we can approximate the posterior distribution, and identify a lower bound of $p(\mathbf{y} \mid y_0)$. In KVAE the posterior distribution is approximated as

$$q(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y}) = q_\phi(\mathbf{x} \mid y_0, \mathbf{y})p_\gamma(\mathbf{z} \mid \mathbf{x}), \quad (2)$$

where $q_\phi(\mathbf{x} \mid y_0, \mathbf{y}) = \prod_{t=1}^T q_\phi(x_t \mid y_0, y_t)$ is parameterized using the inference network, i.e.

$$q_\phi(x_t \mid y_0, y_t) = \mathcal{N}(x_t \mid \mu_t^{\text{enc}}, \Sigma_t^{\text{enc}}). \quad (3)$$

If we rewrite the true posterior distribution

$$p(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y}) = \frac{p(y_0, \mathbf{y}, \mathbf{x}, \mathbf{z})}{p(y_0, \mathbf{y})}, \quad (4)$$

and derive the full distribution model

$$p(y_0, \mathbf{y}, \mathbf{x}, \mathbf{z}) = p(y_0)p_\theta(\mathbf{y} \mid y_0, \mathbf{x})p_\gamma(\mathbf{x}, \mathbf{z}), \quad (5)$$

the true posterior distribution is equivalent to

$$p(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y}) = \frac{p(y_0)p_\theta(\mathbf{y} \mid y_0, \mathbf{x})p_\gamma(\mathbf{x}, \mathbf{z})}{p(y_0, \mathbf{y})} \quad (6)$$

$$= \frac{p_\theta(\mathbf{y} \mid y_0, \mathbf{x})p_\gamma(\mathbf{x}, \mathbf{z})}{p(\mathbf{y} \mid y_0)}. \quad (7)$$

Next, from the KL divergence between the true posterior distribution and our approximate posterior distribution

$$D_{\text{KL}}(q(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y}) \parallel p(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y})) \geq 0, \quad (8)$$

we have that

$$D_{\text{KL}}(q \parallel p) = \mathbb{E}_{q(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y})} \left[\log \frac{q(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y})}{p(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y})} \right] \quad (9)$$

$$= \mathbb{E}_{q(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y})} \left[\log \frac{q_\phi(\mathbf{x} \mid y_0, \mathbf{y})p_\gamma(\mathbf{z} \mid \mathbf{x})p(\mathbf{y} \mid y_0)}{p_\theta(\mathbf{y} \mid y_0, \mathbf{x})p_\gamma(\mathbf{x}, \mathbf{z})} \right] \quad (10)$$

$$= \log p(\mathbf{y} \mid y_0) - \mathbb{E}_{q(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y})} \left[\log \frac{p_\theta(\mathbf{y} \mid y_0, \mathbf{x})p_\gamma(\mathbf{x}, \mathbf{z})}{q_\phi(\mathbf{x} \mid y_0, \mathbf{y})p_\gamma(\mathbf{z} \mid \mathbf{x})} \right] \geq 0 \quad (11)$$

Finally, by moving the expectation to the right-hand side of the inequality, a tractable lower bound of the likelihood is identified

$$\log p(\mathbf{y} \mid y_0) \geq \mathbb{E}_{q(\mathbf{x}, \mathbf{z} \mid y_0, \mathbf{y})} \left[\log \frac{p_\theta(\mathbf{y} \mid y_0, \mathbf{x}) p_\gamma(\mathbf{x}, \mathbf{z})}{q_\phi(\mathbf{x} \mid y_0, \mathbf{y}) p_\gamma(\mathbf{z} \mid \mathbf{x})} \right]. \quad (12)$$

Model architecture

Encoder (129k & 84k parameters): The inference network and the spatial feature extraction share a similar network architecture. We downsample the data using a stack of convolutional layers, where we extract spatial features at each resolution. The network downsamples the data four times using CNNs with filters [32, 32, 32, 16] and then flattens and feeds the features into a dense network. We approximate the posterior distribution by estimating the mean and covariance of x_t .

Decoder (129k parameters): For the generative network, we use attention gates to focus the temporal changes on the spatial features of the reference image at each resolution, followed by an upsampling CNN. The upsampling uses the same number of resolution layers and filters per level as the downsampling. At the output level, we apply a Gaussian filter (with $\sigma_G = 2$ in the ACDC model and $\sigma_G = 4$ in the EchoNet-Dynamic model) after the last convolutional layer. To enforce diffeomorphic estimates of φ_t , we consider the output as the stationary velocity field v_t and compute the transformation numerically using four scaling and squaring layers.

LG-SSM (976 parameters): We design the LG-SSM using eight dimensions for x_t ($p = 8$) and 16 for the state-variable z_t ($q = 16$). We estimate the full matrices A, C , the initial mean μ_0 , and the lower triangular matrices of the covariances R, Q, Σ_0 .

Training procedure: For training purposes, we transform the reference image y_0 using the estimated spatial transformation to compute the likelihood $p_\theta(y_t \mid y_0, \varphi_t)$. For the ACDC experiment, we use a local cross-correlation distribution as likelihood and a Gaussian distribution in the EchoNet-dynamic experiment. We optimize the network using Adam optimizer with a learning rate 5×10^{-4} in both the offline and online scenarios. During offline training, we used a batch size of 4 and trained the ACDC model for 500 epochs and the EchoNet-Dynamic model for 50 epochs.